The Implementation of the U.S. High Performance Computing and Communications Program

Christian Bischof
Math. and Comp. Sci. Div.
Argonne National Laboratory
Argonne, IL 60439

Horst D. Simon
Comp. Sciences Corp.
NASA Ames Research Center
Moffett Field, CA 94035

Summary

This paper gives an overview of the federal High Performance Computing and Communication (HPCC) Program. This initiative involves the Defense Advanced Projects Agency (DARPA), the Department of Energy (DOE), the National Aeronautics and Space Administration (NASA), and the National Science Foundation (NSF). Its goal is to advance U.S. science and industrial competitiveness by supporting the development of teraops high performance systems, the software and algorithms required to render those massively parallel systems efficiently usable, and the Gigabit networks to make those machines accessible, as well as to broaden the base of computationally literate scientists. We outline the HPCC effort and the rationale underlying it, and give an overview of some of the "Grand Challenge" work under way at Argonne's Mathematics and Computer Science Division, as well as the Grand Challenges addressed by NASA's Computational Aerosciences program.

1 The High Performance Computing and Communications Program

The Federal High-Performance Computing and Communications (HPCC) Program was initiated in the President's Fiscal Year (FY) 1992 budget.

The goals of High Performance Computing and Communications Program are to threefold [8]

- Extend U.S. technological leadership in high-performance computing and communications.
- Provide wide dissemination and application of the technologies, both to speed the pace of innovation and to serve the national economy, national security, education, and the global environment.
- Spur gains in U.S. productivity and industrial competitiveness by making high-performance computing and networking techniques an integral part of the design and production process.

The High Performance Computing and Communications program consists of four components:

High Performance Computing Systems: The development of the technology for scalable computing systems in the teraflops range.

Advanced Software Technology and Algorithms: The development of generic software technology and algorithms, as well as systems targeted to Grand-Challenge applications.

National Research and Education Networks: The development of gigabit backbones to provide distributed computing capability to research and educational institutions.

Basic Research and Human Resources: Support mainly for multidisciplinary long-term research and activities to significantly increase the pool of trained personnel.

The High Performance Computing and Communications Program is driven by so-called "Grand Challenges." A Grand Challenge is defined as a fundamental problem in science and engineering, with potentially broad economic, political, and/or scientific impact, that could be advanced by applying High Performance Computing resources. The advancement on the problem should contribute to the productivity, economy, and international competitiveness of the United States, and substantial progress should be expected within the time frame of the HPCC program. The 1993 FCCSET Report lists areas as diverse as magnetic recording technology, rational drug design, high-speed civil transports, catalysis for chemical reactions, fuel combustion, ocean modeling, ozone depletion, digital anatomy, air pollution, design of protein structures, and Venus imaging.

In FY 1992, a HPCC budget of $654.8 million was enacted; and for FY 1993, a budget of $803 million is proposed. About 22 percent of the total funding will be allocated to the High Performance Computing Systems component, 43 percent to the Advanced Software Technology and Algorithms, about 15 percent to the National Research and Education Network, and the remaining 20 percent to Basic Research and Human Resources.

The federal agencies leading the High Performance Computing and Communications program effort, as well as their FY '92 budget and proposed FY '93 budget for the High Performance Computing and Communications program are as follows (in millions of dollars).

Agency	FY 1992	FY 1993
Defense Advanced Research Projects Agency (DARPA)	232.2	275.0
National Science Foundation (NSF)	200.9	261.9
Department of Energy (DOE)	92.3	109.1
National Aeronautics and Space Administration (NASA)	71.2	89.1

2 The Response to a Perceived "Japanese Threat"

In order to justify the High Performance Computing and Communications initiative, the following rationale was put forward [14]:

U.S. leadership in HPC is threatened by Japanese Companies: High-performance computing (HPC) represents the leading edge in information technology and is the part of the information industry where change is occurring most rapidly. The U.S. industry consists largely of "niche players" which cannot amortize the substantial development costs (now well exceeding $100 millions) over a broad product range. As an example, the Gartner report lists Cray Research, the principal U.S. supercomputer vendor, which has annual revenues of $1 billion. By contrast, the three largest Japanese computer companies (Fujitsu, Hitachi, and NEC), have annual revenues ranging from $17 billion to $45 billion [14].

HPC applications are critical: Not only is HPC at the frontier of computing, but computational science, as delivered through HPC applications, is at the frontiers of science, engineering, and related endeavors. HPC offers a competitive advantage, especially in research and development, where it allows more aggressive product goals, shorter time to market, and higher product quality. For example, supercomputers were used for structural design in the Ford Taurus, which reduced the amount of crash-testing necessary. However, Nissan has more supercomputing power than the three U.S. automakers combined.

U.S. government support is essential: U.S. government support is essential for computing technologies in their infancy, especially in light of increasing Asian and European research and development activities. U.S. government assistance in technology diffusion is also imperative, given the close working relationships between government and industry in Japan and (to a lesser extent) in Europe.

The taxpayer's money is well spent: The Gartner Report compared two scenarios. In the first (scenario A), no high-performance computing effort is enacted. In the second (scenario B), $1.9 billion are spent over a five-year period. The following conclusions were made:

Scenario A: The installed supercomputing processing power (measured in peak Mflops) will increase more than 125-fold over the next decade. Of the 176 million peak Mflops in 2000, 90% will be provided by parallel computing systems. Average supercomputer price/performance will increase by a factor of 25 over the next decade, largely as a result of parallel systems.

Scenario B: The installed supercomputing processing power (measured in peak Mflops) will increase more than 300-fold over the next decade. Of the 440 Million peak Mflops in 2000, 96% will be provided by parallel computing systems. Average supercomputer price/performance will increase by a factor of 55 over the next decade.

This report predicted that the cumulative revenue of supercomputing vendors for 1990-2000 would be $10.4 billion greater in Scenario B, and that the industrial applications of this program would increase the U.S. gross national product by $172 to $502 billion over the next decade.

3 Grand Challenge Applications at Argonne

The Department of Energy has long been involved in high-performance computing [2]. In the forties and fifties DOE supported physical research associated with national security. Because of the difficulty, if not the impossibility, of making the pertinent physical measurements, this effort was perhaps the first computational "grand challenge." The DOE Applied Mathematical Sciences program was initiated at the suggestion of John von Neumann to enhance understanding of the use of digital computers in nuclear applications. DOE national laboratories have worked closely with U.S. supercomputer vendors to assist in the design and development of these HPCC systems, to create software environments for use on HPCC systems for both applications and system software components, and to design and implement high-speed local- and wide-area network systems and protocols to provide access to HPCC systems as well as other DOE research facilities.

The DOE HPCC program will build on this experience and work closely with vendors to assist in the design and development of massively parallel systems, in the implementation of applications and system software for these systems, and in the parallel systems access by researchers over the DOE ESNet backbone of the National Research and Education Network. DOE will also build on its current mathematics and computational sciences fellowship and industry exchange programs at the DOE labs to enhance the U.S. human resource base in HPCC technology areas.

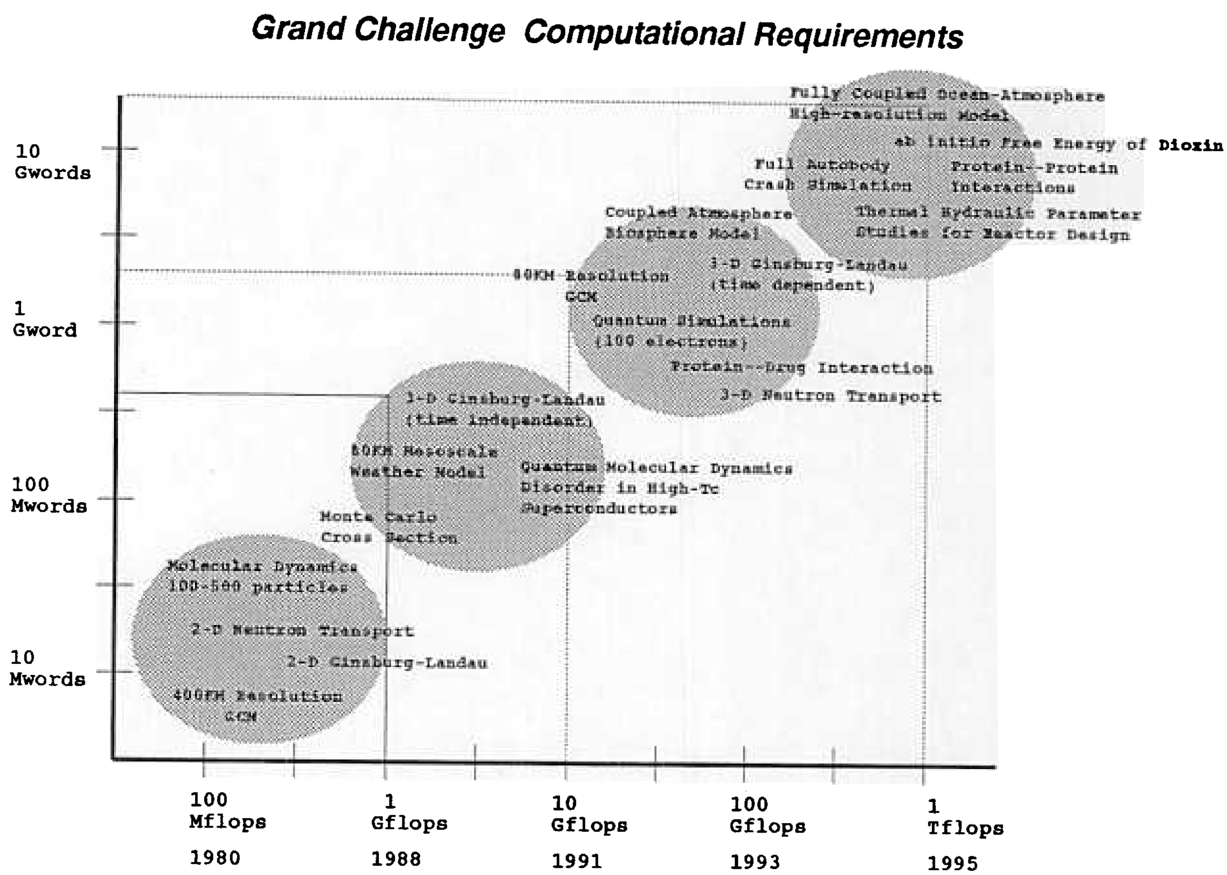

Figure 1: DOE Grand Challenges

Some of the DOE Grand Challenges as well as their computational requirements and the overall time frame are shown in Figure 1.

Solving Grand Challenges requires unprecedented computing power, and to this end Argonne was one of the leading forces in the formation of the Concurrent Supercomputing Consortium (CSC) which operates the Intel Touchstone DELTA System. In addition to Argonne, the institutions that contributed major funding to CSC are the California Institute of Technology, the Jet Propulsion Laboratory, the Center for Research on Parallel Computation (a National Science Foundation Science and Technology Center), DARPA, Intel's Supercomputing Systems Division, NASA, NSF, and Pacific Northwest Laboratory.

The "DELTA" is an impressive computing device. It houses 510 numeric nodes (i860 chips with 16 MBytes of storage each), configured in a 15×34 mesh, 32 I/O nodes (i860 chips with 8 Mbytes of storage and SCSI port), 28 service nodes, and 3 gateway nodes with an Ethernet board. Its peak performance is 30.6 double precision Gflops. It has 8.2 Gbytes RAM, 90 Gbytes of on-line disk storage, and 20 tape drives for archival storage.

We here present some selected DELTA applications in which researchers from Argonne's Mathematics and Computer Science Division collaborated with engineers and scientists on the simulation of atmospheric circulation, the modeling of superconductivity, and the analysis of piezoelectronic crystals.

3.1 Global Climate Modeling

Two conventional numerical methods exist for simulating the general circulation of the atmosphere: the finite difference method and the spectral method. The former suffers the so-called pole problem; that is, the use of a spherical coordinate system in this class of method causes the curvature terms of the momentum equations to be singular at the poles; thus, an artificial filter is needed near the poles. The latter method is expensive in terms of computational time for fine-grid resolution. More important, it is not well suited for parallel supercomputers because it requires global communication of data [9].

The icosahedral method [6] overcomes these deficiencies. It maps the globe onto a mesh of 20 equilateral triangular faces. A triangulation based on some geodesic arguments is then performed on each face, forming an icosahedral grid on the globe. The pole problem is avoided by using this grid and by using the formulation of the governing evolution equations in a three-dimensional Cartesian coordinate system. Another key feature of the new method is the use of a symmetrization procedure. The traditional way of constructing an icosahedral grid on a sphere gives a large variation in distance between grid points, leading to errors. To minimize such errors, three separate grids were constructed; the final grid points are the average of these constructions. This procedure, together with a bilinear interpolation, ensures at least one more digit of precision compared with previous schemes.

The new icosahedral method has been used to perform Rossby-Haurwitz wave simulations, one of the standard tests used to validate the correctness of numerical methods for the shallow-water equations on the sphere. Initial studies indicate that the new method is consistently faster than the well-known spectral transform method on distributed-memory parallel computers. On the Intel Touchstone DELTA System the (unoptimized) icosahedral method [15, 7] achieves a parallel efficiency of about 70% and a performance of approximately 2 Gflops.

The principal investigators of this effort are Ian Foster, William Gropp, John Michalakes, and Rick Stevens. This work is part of the DOE CHAMMP (Computer Hardware, Advanced Mathematical Modelling, and Model Physics Climate Modelling Program) effort [1].

Much of the success in designing and implementing the icosahedral method is contributed to the use of the high-level programming system PCN [11, 10]. PCN itself was developed at Argonne in collaboration with Professors Mani Chandy and Steven Taylor of the California Institute of Technology. PCN made it possible to explore different grid mappings, examine the data graphically to gain an understanding of issues determining parallel performance, and execute the resulting code on a variety of parallel machines without modification.

Analyzing Piezoelectric Crystals

Piezoelectric crystals are important components in electronic appliances such as computers, cellular phones, and pagers. To be useful, these crystals must resonate in a particular vibrational mode at a specified frequency over a wide range of temperatures.

The crystals are modeled by using finite element integration routines, these elements are evaluated and assembled into a sparse matrix data structure in parallel, and the physical configuration of the crystal under thermal stress is then determined by solving a nonlinear equilibrium problem. The dominant computational aspect here is the iterative solution of a sparse system of linear equations. These systems are solved by using advanced linear algebra techniques, in particular new scalable parallel algorithms and software for the solution of linear systems [18, 17]. Once the crystal configuration at the operating temperature is determined, the vibration modes and mode shapes are calculated by solving a large, sparse, generalized eigenproblem. The new software has been used on the Intel DELTA System to model problems involving over 70,000,000 nonzeros and 350,000 equations at a sustained computational rate of around 1.5 Gflops.

The principal investigators in this research are Mark Jones, Paul Plassmann, and Thomas Canfield at Argonne, and Michael Tang at Motorola.

Modeling of Superconductivity

The high transition temperatures achieved in cuprate superconductors in recent years have sparked a great deal of interest in high-temperature superconductivity (HTSC). However, unless mechanisms are found to increase the critical current densities in these superconducting materials, practical applications of HTSC will remain limited. The combination of high temperatures, extreme anisotropy, and intrinsic pinning leads to an extremely rich variety of possible flux vortex behavior. Predicted phases include vortex liquid, entangled vortex liquid, vortex glass (possibly with hexatic bond-orientational order), and a number of pinning-induced regimes of vortex dynamics. Increased computing power is making it possible to consider numerical simulations as a realistic complement to theoretical and experimental research of flux vortex dynamics and phase transitions in superconducting materials.

When a magnetic field is applied to a superconductor, it does not penetrate the material uniformly. Instead, discrete quanta of flux penetrate the material at so-called vortices. The vortices arrange themselves into a hexagonal configuration, In collaboration with researchers in materials science, Hans Kaper, Man Kwong, Gary Leaf, David Levine, J. Moré, Paul Plassmann, and Stephen Wright in the Mathematics and Computer Science Division are working on the study of vortex structures by means of the stationary Ginzburg-Landau equation and the study of vortex dynamics through elastic-filament models and through the time-dependent Ginzburg-Landau equation [19, 13].

Figure 2 shows the magnitude of the order parameter at the minimum energy configuration for a finite element model of a type-II superconductor. This solution was computed on a 120 x 32

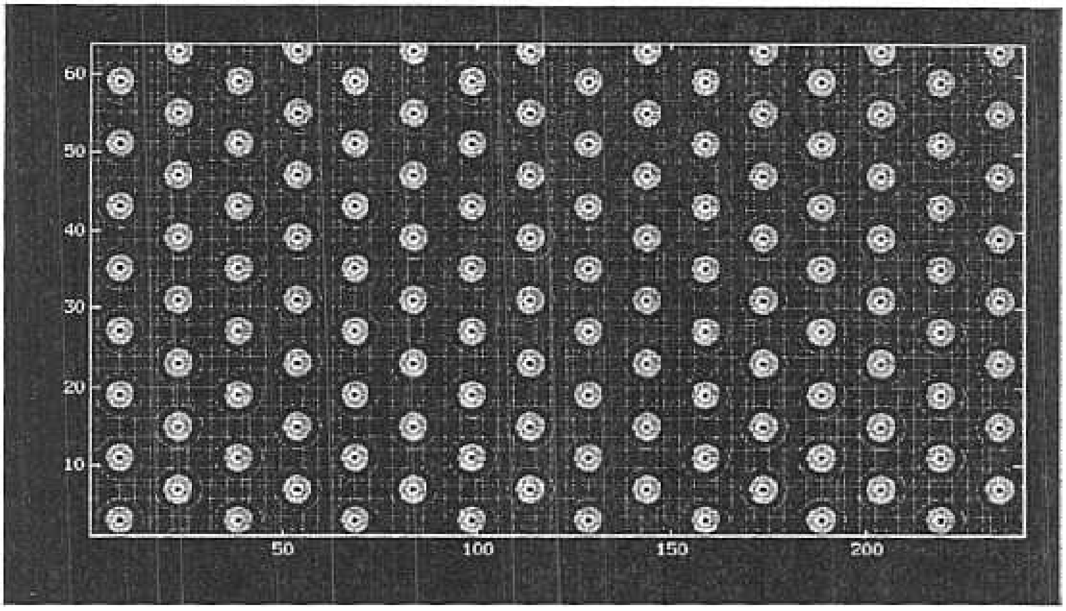

Figure 2: Order Parameter of Minimum Energy Configuration

grid, with quasiperiodic boundary conditions and 32 vortices per cell. In this figure, the cell has been replicated to show the hexagonal pattern of the vortices. This solution was computed using a new, inexact, damped Newton optimization algorithm developed by Paul Plassmann and Stephen Wright.

Figure 3 shows a contour plot of the magnetic field at a minimizer of the homogeneous Ginzburg-Landau problem with 128 vortices in the unit cell. This solution was obtained from a random starting configuration with the inexact, damped Newton method. The data has been replicated one in each direction to demonstrate the periodicity; however, in this case the aspect ratio of the unit cell does not allow for a regular hexagonal tiling by the vortices. What we do see is that the vortices form the hexagonal tiling on smaller "subdomains" with lines where the regular tiling is frustrated between these subdomains. These many vortex solutions give much stronger evidence of the minimum energy vortex configurations than a few vortex solutions with a specially chosen aspect ratio for the unit cell.

4 NASA Grand Challenges

As part of the "Federal High Performance Computing Program", NASA's portion of the "High Performance Computing and Communication Program (HPCCP)" focuses on research and development in areas which show promise to deliver new capabilities to important NASA missions by the late 1990s (for more details see [20]). NASA is planning to

- develop algorithm and architecture testbeds capable of fully utilizing massively parallel concepts and increasing end-to-end performance, and
- develop massively parallel architectures scalable to sustained teraflops performance

These technologies are developed in close interaction with actual Grand Challenge applications. Two NASA Grand Challenges have been chosen as focal points for the NASA HPCCP:

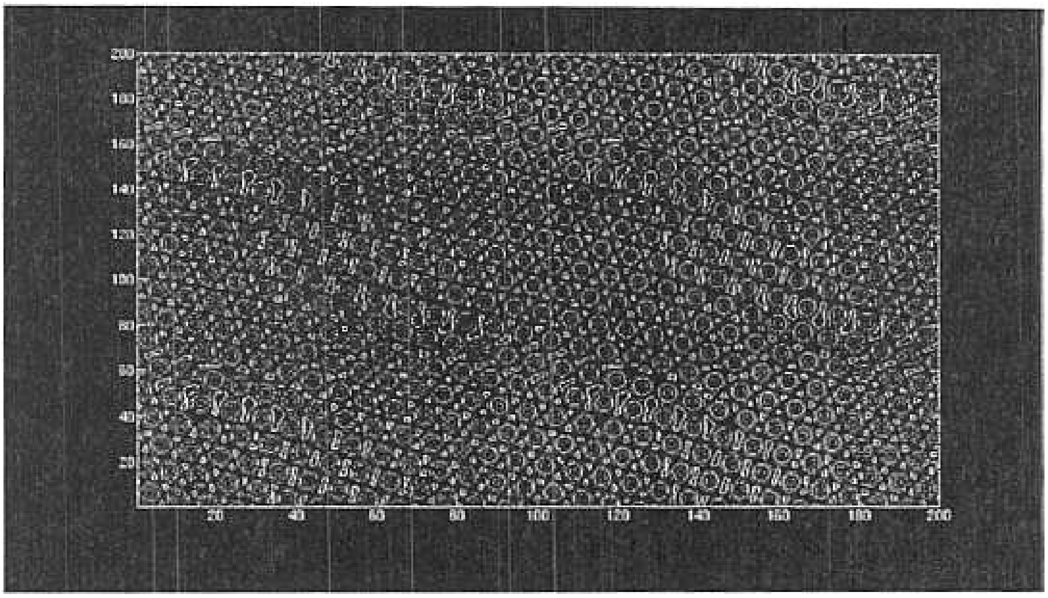

Figure 3: Contour Plot of Magnetic Field

- Computational Aerosciences (CAS) - integrated, multi-disciplinary simulations and design optimizations of aerospace vehicles throughout their mission profiles.
- Earth and Space Sciences (ESS) - multi-disciplinary modeling and monitoring of the earth and its global changes and assessments of their impact on the future environment.

The first Grand Challenge is the focus of the NASA Computational Aerosciences (CAS) program [22, 16]. The CAS goal is to develop the necessary computational technology for the numerical simulation of complete aerospace vehicles for both design optimization and analysis throughout the flight envelope.

Within this program, activities are focused on the development of multi-disciplinary design tools for the high-speed civil transport (HSCT) and high-performance aircraft (HPA). In the high-performance aircraft area, the primary interest is to develop the capability to predict the performance of next generation fighter concepts operating in the most critical portions of their flight regime. To achieve performance levels beyond present generation vehicles, these next generation fighters designs must include higher levels of system integration than can be obtained with present design tools. Towards this goal, aerodynamic, propulsion system, controls, structural, and even acoustic, analysis modules will be integrated into a single software system. The challenges posed by the development and application of such a multi-disciplinary high-performance aircraft analysis tool will be used to illustrate the computational issues in such Grand Challenge computations.

4.1 Grand Challenges of the 1990's (An Example)

Powered-lift aircraft utilize a mix of wing-borne and propulsive lift to achieve vertical or short take-off and landings (V/STOL). With careful design, powered-lift aircraft can also out perform conventional aircraft in other portions of the flight envelope via the use of powered-lift features (e.g., vectoring thrust to achieve super maneuverability). Successful powered-lift aircraft designs are developed from a detailed understanding of the interaction of very complex fluid flows (see Figures in [23]) with all of the major aircraft sub-systems, including the airframe, and propulsion and control systems. Until recently, no computational techniques have been available for the analysis of these complex powered-lift flows [24], and multi-disciplinary interactions [3]. Hence, the design of high-performance powered-lift aircraft has been among the most time-consuming and costly aerospace design activities. As an example, the Harrier was originally conceived in the mid 1950's and is still undergoing significant design studies [12]. Therefore, development of advanced multi-disciplinary analysis tools is being pursued.

A successful computational design tool for high-performance powered-lift aircraft must be able to predict aerodynamic, thermal, and acoustic loads for a vehicle during operations in-ground-effect, transition from jet-borne to wing-borne flight, and in up-and-away flight. Also of key interest is the prediction of engine performance during V/STOL and high-angle-of-attack maneuvers, when inlet flow distortion may degrade thrust or result in engine compressor stall. The V/STOL and transition modes also put severe challenges on the performance of the control system in utilizing the airframe and propulsion systems to retain stable flight.

To model these interactions, at least six computational modules must be integrated:

- Navier-Stokes
- Engine performance
- Structural Heating
- Acoustics
- Control (including pilot model or auto-pilot)
- Aircraft dynamics

Work is presently under way in the Powered-Lift Group of the Applied Computational Fluids Branch at NASA-Ames Research Center towards the Navier-Stokes/Structural Heating/Engine Deck analysis of a Harrier AV-8B in-ground-effect [23]. Work is also underway at NASA- Lewis to develop advanced propulsion system analysis capabilities. Future HPCCP high performance aircraft goals include integrating the aircraft and propulsion analysis tools presently being developed at Ames and Lewis, respectively, into a complete vehicle analysis tool applicable to next generation fighter concepts.

4.2 Surface Modeling and Grid Generation Requirements

A major bottleneck in the application of the described computational design tools will be the development of surface modeling and grid generation software which allows:

1. Surface model definition in less than 1 week
2. Complete grid generation in less than 1 week
3. Design change/regridding of components in less than 1 day
4. Vehicle deformation (e.g., aero-elastic effects) during computation
5. Relative vehicle motion (e.g., landing/take-off) and effector (e.g., flaps and jets) movement during computation

Tasks 1-3 require the development of powerful interactive software tools on workstation platforms, with Task 2 requiring some distributed processing to a super computer (vector or parallel). These requirements are very challenging, but do not necessarily involve parallel computers, and will not be addressed in detail here.

Tasks 4 and 5 must be performed during the numerical simulation on the parallel computer systems. Accommodating vehicle deformation (Task 4) will require that a parametric representation (e.g., NURBS) of the vehicle surface reside on the parallel computer, and that this geometric representation can be manipulated and sampled dynamically without user intervention. It will also be required that new volume grids be created dynamically, using the deforming vehicle geometry as the new boundary condition for the algebraic or PDE (e.g., elliptic) volume grid generator. Accounting for vehicle and effector movement (Task 5) will be best accommodated using an overset grid technology (e.g. [4, 5]). In this case, as the aircraft moves in relationship to the ground (for example) the grids attached to the aircraft and ground will be in relative motion, and new interpolation stencils must be computed at each iteration. This requires that the nearest-point and interpolation features of the overset-grid technology be ported to the parallel computers. Considering that the technology required for Tasks 4-5 is only in the formative stages of development on vector computers, the challenge of fully-developing this software and implementing it in the parallel environment is formidable.

4.3 Flow Simulation (CFD) Requirements

In 1991, a state-of-the-art simulation of the flow about a Harrier operating in-ground effect required approximately 2.8 million points, 20 Mwords of run-time memory, and about 40 hours of CPU time on a Cray Y-MP running at a sustained speed of approximately 160 MFLOPS. Such a computation solves the Navier-Stokes equations for the viscous flow about the Harrier using, in this case, a simple algebraic turbulence model. The grid was the coarsest possible that would still allow most of the important flow features to be resolved. The predicted flow features are in good agreement with flight flow visualization [23].

It is estimated that to obtain "engineering-accuracy" predictions of surface pressures, heat transfer rates, and overall forces, the grid size will have to be increased to a minimum of 5.0 million points. If the unsteady motion of the flow structures is to be resolved, at least 50,000 iterations will also be required. Also, more advanced turbulence modeling must be included. In summary, we anticipate the following minimum requirements in terms of floating point operations for just the external flow simulation element of future Grand Challenge computations:

- 5,000,000 grid points
- 50,000 iterations
- 5,000 operations per point per iteration
- 10^{15} operations per problem

The actual computational speed requirements for such a calculation depend on the mode in which the calculation is carried out. In a proof-of-concept mode such a calculation may be carried out only once as a "heroic effort". If this could be done in 100 to 1000 hours turn-around-time, it would translate into a sustained speed between 3 and 0.3 GFLOPS. Design and automated design modes require a much lower turn-around-time and thus result in much higher requirements for computational speed. The corresponding figures are summarized in Table 1.

Table 1: **Requirements for Flow Simulation**

Solution Mode	Turn-around-Time	Required Performance
Proof-of-concept	1000 – 100 hours	0.3 – 3 GFLOPS
Design	10 – 1 hours	30 – 300 GFLOPS
Automated Design	0.1 – 0.01 hours	3 – 30 TFLOPS

These computational requirements are accompanied by a corresponding increase in memory and storage requirements. Approximately 40 storage locations are required per grid point. If all of the computational zones remain in memory, this translates to a requirement for 200 million words of run-time memory (to date, often a desirable feature for parallel systems). For unsteady flow analysis 100-1000 time steps (at 8 words per point) must be stored. This leads to a requirement of 4-40 gwords of "disk" storage per problem.

If we compare these requirements with the computer resources required to address the "Grand Challenges" of the 1980's (e.g., a 1.0 million point steady Navier-Stokes simulation, on a Cray-2 class machine, of the external flow about an aircraft at cruise) we arrive at Table 2.

Table 2: **Proof-of-Concept Requirements: 1980's vs. 1990's**

	1980's	1990's	Ratio
100 hr. run time	40 MFLOPS	3000 MFLOPS	75
run-time memory	35 Mwords	200 Mwords	6
"disk" storage	8 Mwords	40000 Mwords	5000

We note in particular that a 5000 fold increase in data storage and manipulation capabilities will be required to address CFD Grand Challenges of the 1990's. A single solution file for a time step will have up to 40 Mwords (320 Mbytes) of data. The above discussion assumes that the computation for advancing the solution one time step can be carried out in about 10 seconds. Even though it is not necessary to store a solution file at every time step, these figures show the need for a sustained I/O bandwidth of at least 40 Mbytes/sec. For a more detailed discussion of I/O requirements for parallel CFD see the report by Ryan [21].

4.4 Grand Challenges of the 1990's: Multi-disciplinary Computations

The discussion in the previous section was restricted to prediction of the external flow about an advanced fighter concept. As explained in subsection 4.1 the Grand Challenge computations of the 90's will be multi-disciplinary, combining computational techniques useful in analyzing a number of individual areas such as structures, controls, and acoustics, in addition to the baseline CFD simulations. It is possible in all these areas to derive estimates for the performance requirements. These estimates are given in Table 3 as multiplicative factors of additional requirements over the single-discipline baseline CFD simulation.

It is clear that computational resource requirements can increase rapidly for multi-disciplinary computations. If the corresponding factors for multi-disciplinary V/STOL aircraft design are extracted from Table 3, and combined with the numbers for the baseline external aerodynamics prediction, quickly Gword and near TFLOP requirements arise. The details are given in Table 4 [20].

Table 3: **Increase in memory and CPU Requirements over Baseline CFD Simulation**

Discipline	**Memory increase**	**CPU Time increase**
Structural Dynamics		
modal analysis	× 1	× 2
FEM analysis	× 2	× 2
thermal analysis	× 2	× 2
Propulsion		
inlet/nozzle simulation	× 2	× 2
engine performance deck	× 2	× 2
combustion model, e.g. scramjet	× 4	× 10
turbojet engine (full sim.)	× 10-100	× 10-100
Controls		
control law integration	× 1	× 1
control surface aerodynamics	× 2	× 2
thrust vector control	× 2	× 2
control jets	× 2	× 2
Acoustics	× 10	× 10
Numerical Optimization Design	× 2	× 10-100

Table 4: **Flops and Run-time Memory Requirements for 5 Hour Run.**

	Mwords	**GFLOPS**
Base CFD	200	60
Structural		
thermal analysis	× 2	× 2
Propulsion		
inlet/nozzle simulations	× 2	× 2
engine performance deck	× 2	× 2
Controls		
control law integration	× 1	× 1
thrust vector control	× 2	× 2
Total	2000	600

It should be noted that the factors in Table 3 are based on the assumption that the physical frequencies introduced because of the multi-disciplinary integration can be resolved with the time steps required by the aerodynamics simulation. Additional compute time may be required if the multi-disciplinary system exhibits higher frequency modes which must be resolved.

5 Summary

This paper gave an overview of the U.S. High Performance Computing and Communications Program as well as some of the Grand Challenges pursued by the Department of Energy and the National Aeronautics and Space Administration. The goals of this program are ambitious, and the integrated approach favored in its pursuit is a marked departure from the previously favored "single-discipline" funding. In this fashion, this initiative is likely to have a marked impact on the future of "computational science" in the United States, and in the world.

Acknowledgments. We are indebted to William R. Van Dalsem of NASA Ames Research Center. The material in Section 4 has been originally developed by him and is reprinted from [22]. We would also like to thank Thomas Morgan, Gail Pieper, Paul Plassmann, and Stephen Wright for their contributions to this summary report.

This work was supported by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U.S. Department of Energy, under Contract W–31–109–Eng–38, and by the Applied Research Branch, Numerical Aerodynamic Simulation (NAS) Systems Division, NASA Ames Research Center through NASA Contract NAS 2-12961.

References

[1] Building and advanced climate model: Program plan for the CHAMMP climate modelling program. Technical Report DOE/ER-0479T, Department of Energy, 1990.

[2] The doe program component of the federal high performance computing and communications program. Technical Report DOE/ER-0489P, Department of Energy, 1991.

[3] P. A. Abeloff, W. R. Van Dalsem, and F. C. Dougherty. Thermal interaction between an impinging hot jet and a conducting solid surface. 1990. AIAA Paper 90-0299.

[4] J. A. Benek, P. G. Buning, and J. L. Steger. A 3-d chimera grid embedding technique. 1985. AIAA Paper 85-1523.

[5] J. A. Benek, T. L. Donegan, and N. E. Suhs. Extended chimera grid embedding scheme with application to viscous flows. 1987. AIAA Paper 87-1126-CP.

[6] I-L. Chern. A control volume method on an icosahedral grid for numerical integration of the shallow-water equations on the sphere. Technical Report MCS-P214-0291, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[7] I-L. Chern and I. Foster. Design and parallel implementation of two numerical methods for modeling the atmospheric circulation. Technical Report MCS-P264-0991, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[8] Mathematical Committee on Physical, Engineering Engineering Sciences, Federal Coordinating Council for Science, and Technology. Grand challenges 1993: High performance computing and communications, 1992.

[9] I. Foster, W. Gropp, and R. Stevens. The parallel scalability of the spectral transform method. Technical Report MCS-P215-0291, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[10] I. Foster, R. Olson, and S. Tuecke. Productive parallel programming: The PCN approach. Technical Report MCS-P295-0392, Mathematics and Computer Science Division, Argonne National Laboratory, 1992.

[11] I. Foster and S. Tuecke. Parallel programming with PCN. Technical Report ANL-91/34, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[12] J. W. Fozard. *The Jet V/STOL Harrier - An Evolutionary Revolution in Tactical AirPower.* AIAA Professional Study Series, 1978.

[13] J. Garner, M. Spanbauer, R. Benedek, K. J. Strandburg, S. Wright, and P. Plassmann. Critical fields of Josephson-coupled superconducting multilayers. Technical Report MCS-P281-1291, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[14] Gartner Group. High performance computing and communications: Investment in American competitiveness, 1991.

[15] S. J. Hammond. A multilingual parallelization of the icosahedral-hexagonal grid system for integration on the sphere. Technical Report MCS-P234-0491, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[16] T.L. Holst, M. D. Salas, and R. W. Claus. The NASA computational aerosciences program - toward teraflop computing. In *Proceedings of the 30th Aerospace Sciences Meeting, Reno, NV*, Washington, D.C., 1992. American Institute of Aeronautics and Astronautics. AIAA Paper 92-0558.

[17] M. T. Jones and P. E. Plassmann. Fortran subroutines to compute improved incomplete Cholesky factorizations. Technical Report MCS-P221-0191, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[18] M. T. Jones and P. E. Plassmann. A parallel graph coloring heuristic. Technical Report MCS-P246-0691, Mathematics and Computer Science Division, Argonne National Laboratory, 1991.

[19] Hans Kaper. Mathematical models of superconductivity. Technical Report MCS-TM-146, Mathematics and Computer Science Division, Argonne National Laboratory, 1990.

[20] National Aeronautics and Space Administration. *Project Plan for the Computational Aerosciences Project of the NASA High Performance Computing and Communications Program (NASA HPCCP/CAS)*, May 1991.

[21 James S. Ryan. Concurrent File System (CFS) I/O for CFD Applications. August 1991.

[22] Horst D. Simon, William Van Dalsem, and Leonardo Dagum. Parallel CFD: Current Status and Future Requirements. In Horst D. Simon, editor, *Parallel CFD - Implementations and Results Using Parallel Computers*, pages 1 – 28. MIT Press, Cambridge, Mass., 1992.

[23] M. H. Smith, K. Chawla, and W. R. Van Dalsem. Numerical simulation of a complete STOVL aircraft in ground effect. 1991. AIAA Paper 91-3293.

[24] W. R. Van Dalsem, K. Chawla, M. H. Smith, and P. A. Abeloff. Numerical simulation of powered lift flows. In *Proceedings of the International Powered-Lift Conference*, London, England, August 1990. The Royal Aeronautical Society.